

GreenCore: Green Computing for the Future of DeFi

The GreenCore Team
xyz@supercompiler.xyz

Abstract

The blockchain paradigm has demonstrated its utility to unleash many new classes of decentralised applications, with DeFi as a prominent example. There have been a number of ongoing projects to address a key challenge in DeFi: the high transaction fees. We introduce a new solution, GreenCore, that alleviates this problem based on optimization theories of computer programs. GreenCore analyzes the deployed bytecodes of a smart contract and automatically optimizes its gas cost on every code path. We discuss its design, implementation and the opportunities it provides for the future. We further develop GreenSwap, by optimizing the popular Uniswap protocol, and find that it significantly reduces transaction fees on Ethereum by over 30% on average compared to Uniswap.

1 Introduction

Blockchain technology has materialized the idea of building a world computer: a composable, open, permissionless state machine that runs trust-minimized code without borders. While we are still at an early stage of this revolution, DeFi (Decentralized Finance) has emerged as one of the most successful applications atop this new computer. Unlike traditional finance, DeFi allows anyone to borrow, invest, trade, issue and own digital assets without a centralised intermediary. This not only makes financing more efficient, but more importantly creates a transparent and more resilient financial system as compared to traditional banks.

In early 2021, the total market capitalization of DeFi has exceeded \$100 billion. For many users, a pain in using DeFi is its high transaction fee, also known as “gas” fee. A simple exchange of two cryptocurrencies on Ethereum, regardless of the exchange amount, often costs more than \$10, and can sometimes go up to over \$1000. This problem has inspired a number of notable innovations to reduce transaction fees, such as Binance Smart Chain, Solana, as well as various Layer-2 solutions. However, none of them is perfect: they all rely on exploring different tradeoffs among three important elements—decentralization, security and scalability.

In this paper, we introduce a new solution in the design space—GreenCore—to alleviate this problem. GreenCore is based on optimization theories of computer programs. Observing that transaction fee is a product of two multiplication factors: gas price and gas cost of the transaction. While the former can fluctuate depending on the situation of network congestion, the latter is determined by the smart contracts,

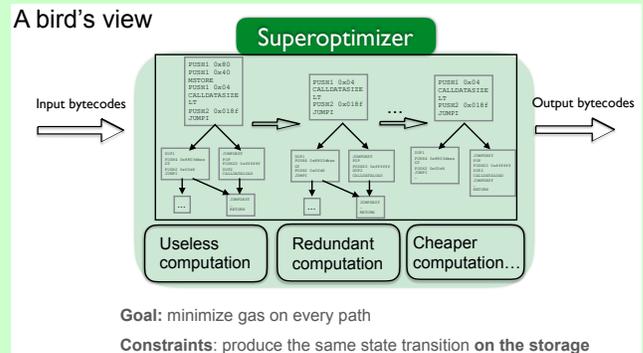


Figure 1. GreenCore.

i.e., the code executed on the world computer during the transaction. GreenCore automatically optimizes the code towards reducing its gas cost, which consequently lowers the transaction fee by a multiplying factor. The more optimized the code is, the cheaper transaction fee we get.

Figure 1 depicts a bird’s view of GreenCore, which takes a sequence of bytecodes as input, optimizes them with a Superoptimizer, and outputs another sequence of optimized bytecodes. Every smart contract is a sequence of bytecodes when deployed to the world computer. The output code from GreenCore has the same format as the input code and can be deployed with no changes to the world computer. Moreover, the output code is functionally equivalent to the input code: under all scenarios allowed by the world computer, the output code always produces the same state transition as the input code (note that the computer is a state machine).

GreenCore’s cornerstone is an automated Superoptimizer driven by *Maximal Concolic Execution (MCE)*, a code verification technique developed by the programming language research community. MCE provides the ability to systematically explore every code path, encode path conditions with mathematical formulas, and check path invariants with SMT (Satisfiability Modulo Theories) solvers. The Superoptimizer can then identify optimization opportunities, such as useless computation, redundant computation, cheaper computation and so on, to minimize gas cost on every code path. Finally, the optimization is realized by iteratively transforming an intermediate graph representation of the input code while respecting the path invariants.

Powered by GreenCore, we have developed GreenSwap, a highly optimized DEX based on Uniswap—an automated market maker (AMM) on Ethereum. Uniswap is among the most popular DeFi protocol for swapping ERC20 tokens, and

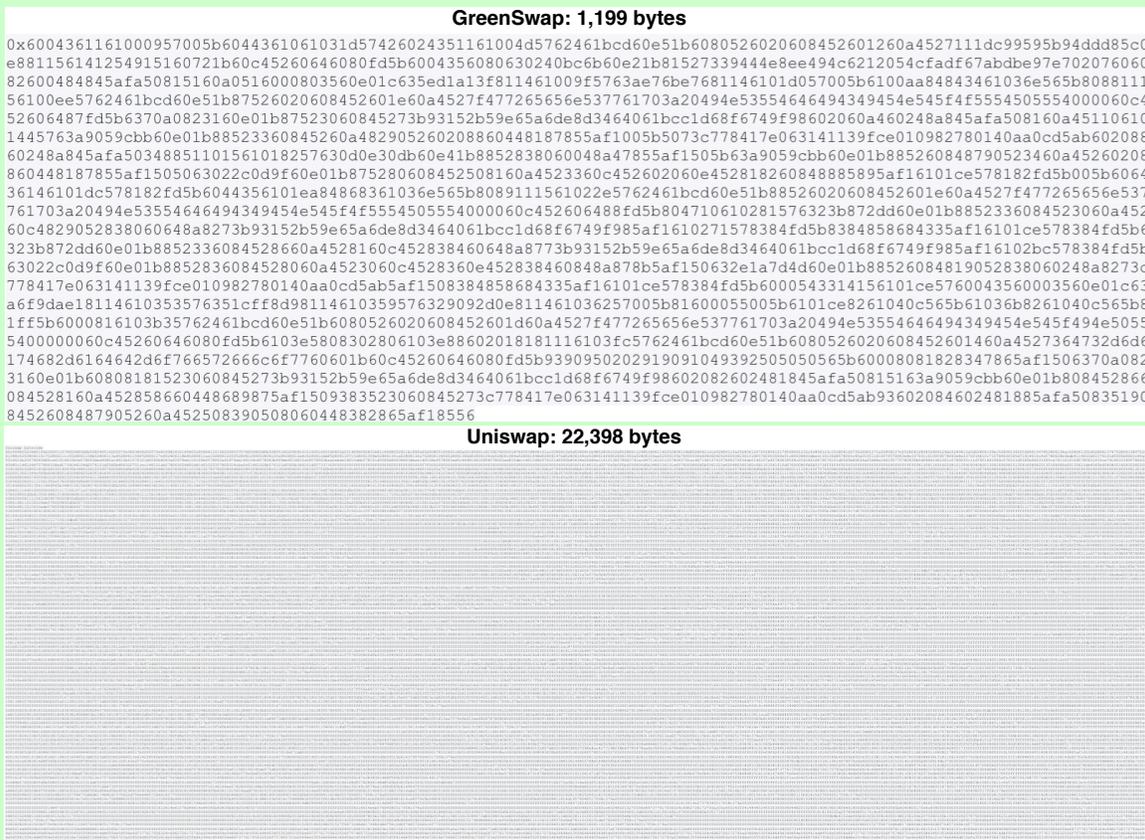


Figure 2. Deployed bytecodes of GreenSwap and Uniswap.

its code has been executed over 50 million times by millions of users through swap transactions. A swap transaction on Uniswap takes about 150K gas cost and \$20 gas fee on average. Compared to Uniswap, GreenSwap reduces the gas cost by as much as 55% (over 30% on average), saving as much as \$10 fee per transaction. Figure 2 shows the bytecodes of GreenSwap and Uniswap (UniswapV2Router02). The size of GreenSwap is almost 18X smaller than Uniswap. This result indicates that there are abundant optimization opportunities in Uniswap identified by GreenCore. Note that the gas cost of a transaction is determined by the executed bytecodes on the code path during the transaction, not the size of the whole smart contract bytecodes. In our tests on Ethereum Ropsten, GreenSwap reduces the gas cost of Uniswap to only around 70K, more than 50% reduction on average. Once it is widely deployed and used by DeFi users, GreenSwap will save a significant amount of transaction fees.

We envision that GreenCore can serve as a foundation to make DeFi more efficient and ESG-friendly (Environmental, Social, Governance). GreenCore has also a large potential to be used for ensuring security of smart contracts. When used for auditing purposes, it can eliminate the gap of compiler bugs, which may introduce new security vulnerabilities in the compiled bytecodes unseen from the source code.

2 The GreenCore Technology

We use a simple flipper contract written in Solidity and compiled for the EVM to illustrate the GreenCore technology.

2.1 EVM

The Ethereum Virtual Machine (EVM) is a stack machine that executes all the bytecodes deployed on Ethereum. Most bytecodes compute data on the stack, or read/write data on a temporal memory or the persistent storage which is expensive. There are a few control bytecodes that determine what to execute next, such as JUMP/JUMPI/CALL/STOP/RETURN/REVERT. The first step of GreenCore is to transform the input sequence of bytecodes into an intermediate graph representation that respects the execution flow determined by the control bytecodes. Figure 3(a) shows the source code of the flipper contract. Figure 3(b) shows the bytecodes generated by the Solidity compiler with optimization turned on. Figure 3(c) shows a graph constructed from the bytecodes. Each node in the graph contains a subsequence of the input, starts with a tag (a jump target excepts the entry node) and ends with a control bytecode. Each edge in the graph represents the execution flow from a control bytecode to a jump target. Note that the graph may not be straightforward to construct, when a jump

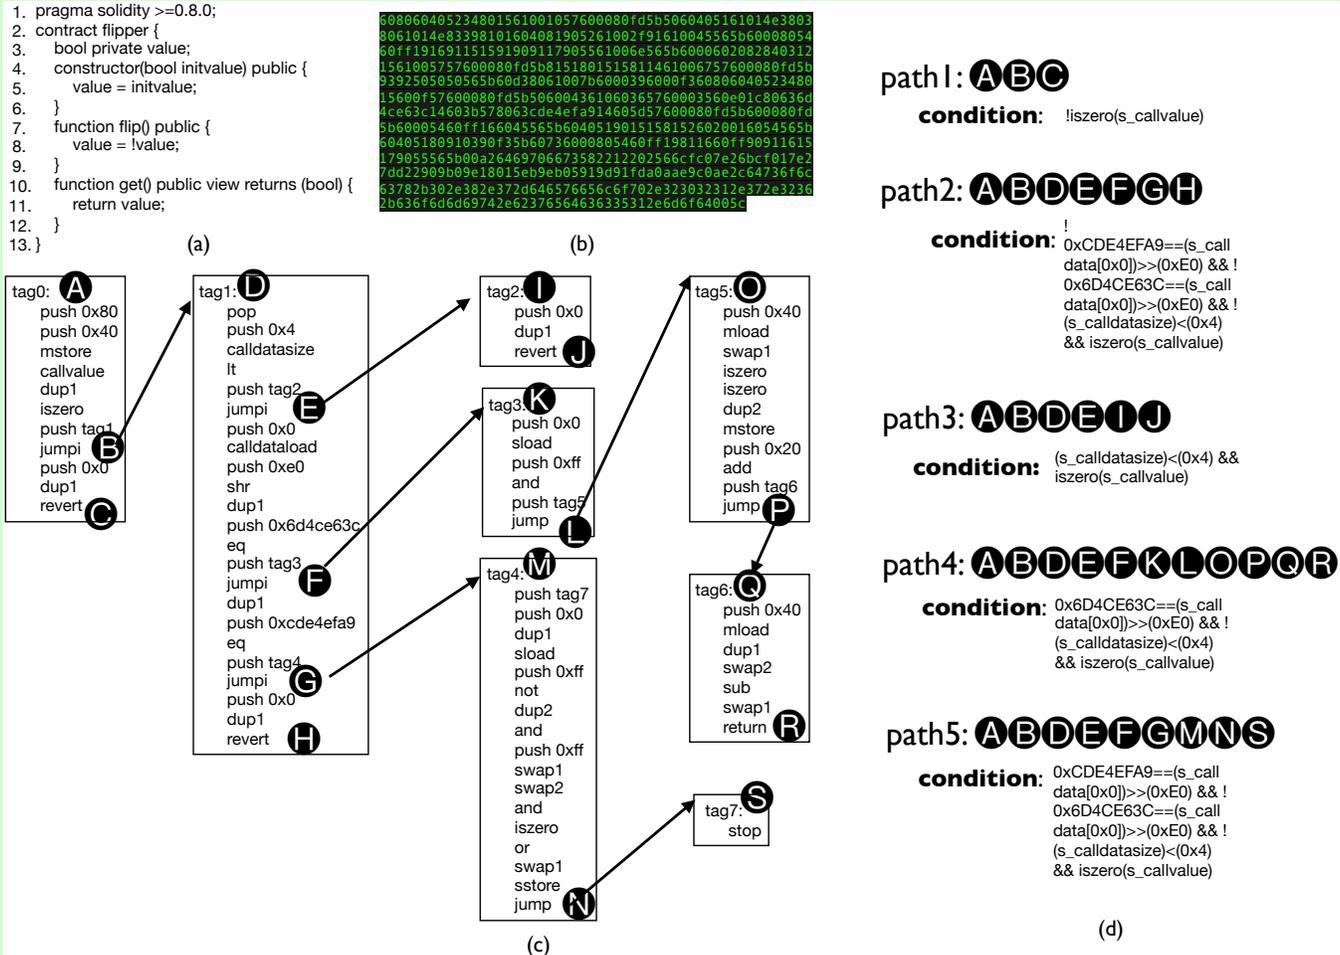


Figure 3. An illustration of GreenCore: (a) a smart contract written in Solidity, (b) bytecodes produced by Solidity compiler with optimization, (c) graph representation of the bytecodes, (d) code paths and the corresponding path conditions.

target depends on computed data. This problem is addressed by Maximal Concolic Execution.

2.2 Maximal Concolic Execution

Maximal Concolic Execution (MCE) is a technique that uses both concrete values and symbolic values to efficiently test and verify code. It executes the input bytecodes by means of interpretation, with the goal to explore all code paths. Every code path starts from the input entry and ends with a terminating bytecode STOP/RETURN/REVERT. Whenever data is unknown, such as input data from the contract call or data returned from an external call, it will introduce a symbolic value to denote the unknown data, and the interpretation will continue using the symbolic values. Upon a control bytecode, MCE will determine the possible jump targets based on the path condition, which may contain symbolic values. The process is repeated until no new code paths can be found. Figure 3(d) shows the five paths explored by MCE for our example. Each path is associated with a path condition

denoting that under what condition the path is valid. For example, path1 has condition !iszero(s_callvalue), meaning that this path will be executed when the contract is called with value. In more complex cases, invalid paths will be pruned by solving the corresponding path conditions with SMT solvers, which are often highly efficient.

2.3 Superoptimizer

The Superoptimizer identifies optimization opportunities based on the explored paths, and performs the optimizations by transforming the graph representation. The optimization opportunities can be characterized with respect to removing useless and redundant computation, and replacing expensive computations with equivalent but cheaper computation. Figure 4 shows a large number of optimizations identified in our example. For example, under tag0 the first three bytecodes push 0x80 push 0x40 mstore can be removed because the data stored by mstore is never used. The dup1 under tag0 and pop under tag1 are redundant and can be

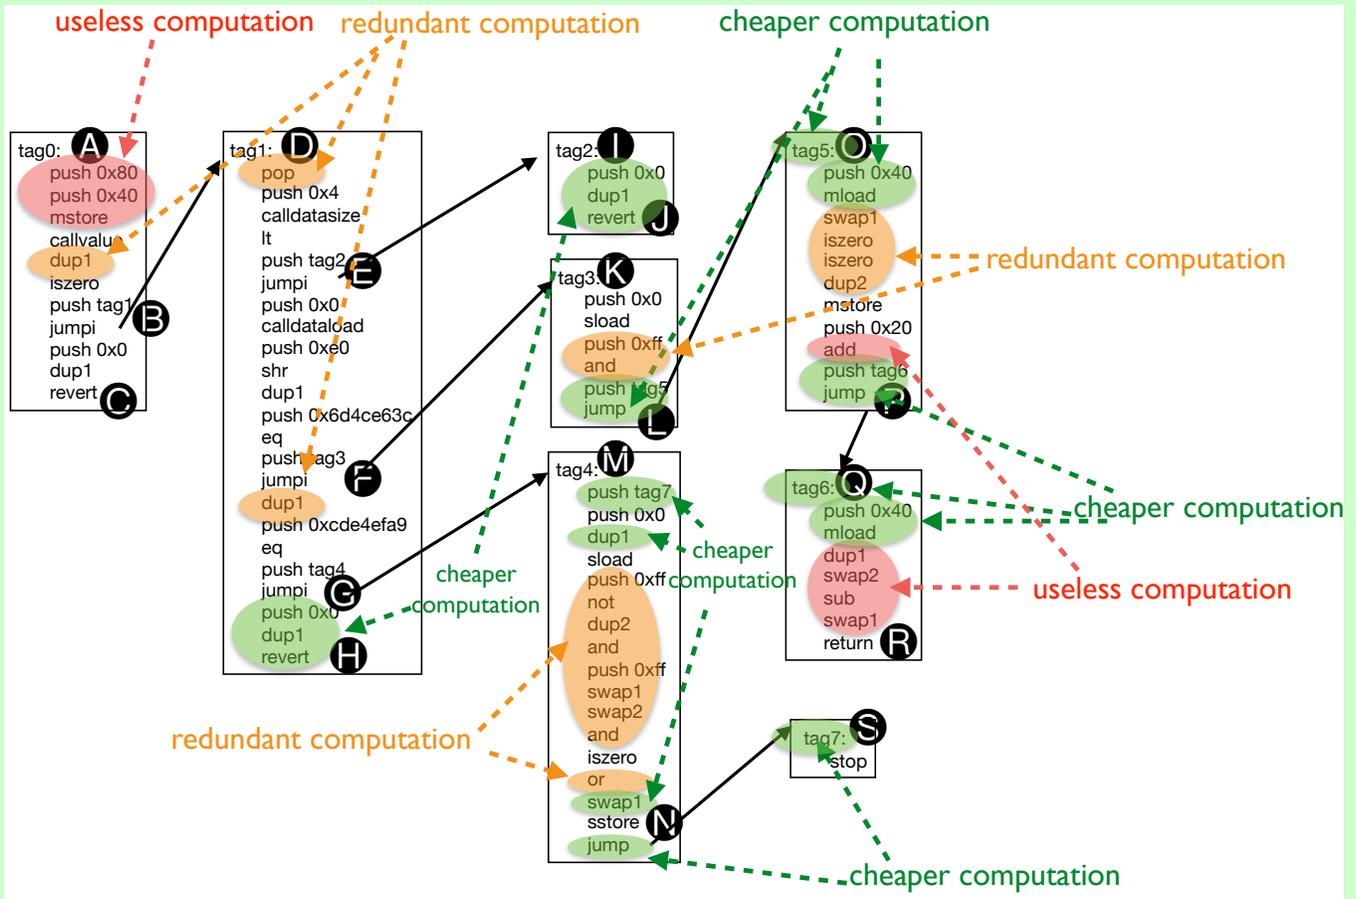


Figure 4. GreenCore Superoptimizer.

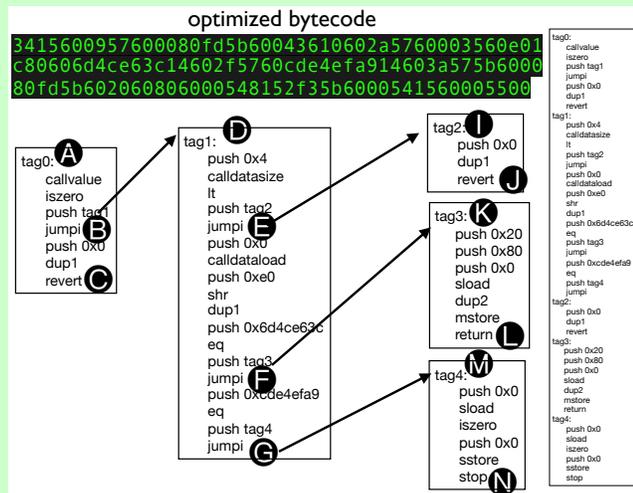


Figure 5. GreenCore optimized bytecodes.

both removed, because on path1 the value by dup1 is not used, and on all the remaining four paths the value is dropped

by pop. The jump under tag4 can be replaced by stop under tag7, which reduces the cost of a jump. The other cases are similar. After all the optimizations, GreenCore performs a topological sort of the graph (prioritizing hot code paths) and produces the output bytecodes by replacing tags with concrete addresses. Figure 5 shows the optimized bytecodes for our example. Compared to the input bytecodes in Figure 3(b), the size of optimized bytecodes is 4X smaller.

3 Conclusion and Future Directions

We have introduced GreenCore and illustrated the technical approach. We have also developed GreenSwap as one application of GreenCore. There are a number of future directions that are promising to further explore the potentials of GreenCore. In the immediate future, GreenCore can be applied more widely to other DeFi applications besides GreenSwap. GreenCore can be also used to audit and secure smart contracts by automatically verifying code paths at the level of bytecodes. In the longer term, GreenCore can be developed as a foundational tool to optimize and secure future decentralized applications and blockchains beyond Ethereum.